

VIDEO-ON-DEMAND VIDEO SERVER**DISK/MEMORY STREAMING SELECTION METHODOLOGY**

5

CROSS REFERENCE TO CO-PENDING APPLICATIONS

elmo.a7 This application is related to commonly assigned and co-pending U.S. Patent Application Serial No. 33012/263/101, filed _____, and entitled "Video on Demand Transaction Server"; U.S. Patent Application Serial No. 33012/264/101, filed _____, and entitled "Video Server"; U.S. Patent Application Serial No. 33012/287/101, filed _____, and entitled "Set Top Network Protocol"; U.S. Patent Application Serial No. 33012/262/101, filed _____, and entitled "Video On Demand System"; U.S. Patent Application Serial No. 33012/285/101, filed _____, and entitled "Menuing Application for Video On Demand System"; U.S. Patent Application Serial No. 33012/265/101, filed _____, and entitled "Video on Demand Transaction Gateway"; U.S. Patent Application Serial No. 33012/265/101, filed _____, and entitled "Video on Demand Transaction Gateway"; and U.S. Patent Application Serial No. 33012/266/101, filed _____, and entitled "A Web Based Video on Demand Administration Application", all of which are incorporated by reference.

BACKGROUND OF THE INVENTION

1. **Field of the Invention:** The present invention generally relates to digital data transmission of video information and more particularly to the delivery of user selected video information to subscribing users.

5

2. **Description of the Prior Art:** The mass distribution of video programming signals (i.e., television) was originally accomplished primarily by the broadcasting of a very high frequency (i.e., VHF) carrier containing an amplitude modulated video signal and a frequency modulated audio signal. Through the addition of more broadcasting transmitters, a modest number of 10 different programming signals could be simultaneously distributed to a large number of potential users with a modest capital commitment. Using this technique, the capital commitment increases almost linearly with number of different programming channels within the limits of the available spectrum space for separate and independent carriers.

Within a couple of decades, most of the scarce VHF spectrum space had been committed, 15 and increasing demand for additional programming channels resulted in the allocation of spectrum space in the ultra-high frequency (i.e., UHF) region. Whereas virtually all receivers became UHF compatible, as a matter of policy, and UHF channels were assigned to requesters, it was appreciated that there were no economies of scale through the addition of more broadcast programming channels.

20 The cable television channel era was the result, wherein a capital commitment was required to wire each user home within a service area. As a result, about one hundred separate

programming channels became readily available at a lesser cost than providing the same number of channels using conventional broadcast means. Initially, cable television was simply an analog system in which low power modulated carriers were transferred over a coaxial cable rather than being broadcast into the ether at substantially higher power levels. The cost saving was realized
5 because the broadcasting was accomplished at substantially lower power.

With approximately one hundred different programming channels, it is typical to charge cable service user fees in accordance with a hierarchy of programming channels. The least expensive channels tend to be the preexisting broadcast channels and those cable channels supported primarily by advertisers which are intended for the most general audiences. The
10 subscriber fees to access other channels increase as the programming becomes more specialized, advertising revenue becomes less likely to pay the programming costs, and the programming materials tend to have substantial economic value through other distribution channels. The so-called "premium channels" which show current and/or near current movies without advertising are typical of the higher cost programming options.

15 Most commercial cable television providers package the various programming channels into programming channel groups with different prices such that a given user can select a suitable programming package and pay the equivalent fee. Typically, a cable provider box, which couples the user television receiver to the coaxial cable source, is controlled by the cable television provider to give access to a given user to only those channels for which the appropriate subscriber
20 fee has been paid.

The most expensive cable television channels currently available are "pay-per-view" or PPV. With the PPV concept, a given user can subscribe to a given programming channel for a single individual program of up to several hours for a separate subscriber fee. Typically, PPV channels provide sporting events and almost current movies.

5 Perhaps the major disadvantage of the PPV concept as currently implemented, is that the programming is provided in the "broadcast" mode. That means that the programming begins and runs on a predefined schedule. As a result, programming is missed if the user receives a telephone call, for example, during the viewing. Furthermore, it ordinarily requires the user to allocate viewing time to coincide with the predefined schedule. To overcome this disadvantage, many
10 users rent video programs as video cassette recordings (i.e., VCR) from commercial stores which provide such a rental service. This permits the viewer to watch the program in accordance with her/his own schedule, stop the program during interruptions, and replay portions of the program which may not be readily understood. The primary disadvantage of the VCR rental approach is the need to physically go to the rental store to obtain the program and return to the rental store to
15 return the recording.

With the capital commitment for cable television in place, there appear to be substantial new uses for the basic coaxial pathway. Such uses include, telephone, computer modem, facsimile, and video conferencing. To properly coordinate such diverse information transmission activities, attention is being directed to digital transmission schemes which provide for easier
20 management of the distribution resources. U.S. Patent No. 5,570,355, issued to Dail et al., discusses the handling of a number of diverse information transmissions within a single system.

U.S. Patent No. 5,673,265, issued to Gupta et al., U.S. Patent No. 5,754,773, issued to Ozden et al., and U.S. Patent No. 5,799,017, issued to Gupta et al., all discuss multi-media distribution systems. U.S. Patent No. 5,555,244, issued to Gupta et al., is directed to multimedia distribution to residential users.

5 The digitization of video results in a great deal of data which must be transferred at a high rate to yield acceptable performance and resolution. By current standards, 3mbits/sec. is considered to be a very acceptable rate. Such high data rates require systems which can provide high data rate transmission. U.S. Patent No. 5,724,543, issued to Ozden et al., U.S. Patent No. 5,699,362, issued to Makam, and U.S. Patent No. 5,826,110, issued to Ozden et al., all concern 10 themselves with high data rate retrieval and transmission. U.S. Patent No. 5,675,573, issued to Karol et al., discusses the management of high data rate bandwidths.

In addition to retrieval and transmission of the required high data rates, there is also the need to provide high speed switching for switching as between data sources and destinations. U.S. Patent No. 5,751,704, issued to Kostic et al., and U.S. Patent No. 5,740,176, issued to 15 Gupta et al., discuss high speed digital switching systems.

Whether it is data storage and retrieval, data transmission, or data switching, the fundamental technological problem associated with digital video results from the sheer volume of digitized video data and the tremendous rate at which it must be provided to the ultimate user for satisfactory performance. One technique for reduction of the volume problem is in reducing the 20 resolution (and hence the volume of data) for those applications for which such reduction is acceptable. U.S. Patent No. 5,623,308 and U.S. Patent No. 5,691,768, both issued to Civanlar et

al., directly address the handling of multiple resolution digitized video signals within a single system.

Notwithstanding attempts to reduce the resolutions to the lowest acceptable levels, the total data volume of any commercially useful system will remain high. The most common way to 5 treat extremely high data volumes is through data compression. U.S. Patent No. 5,710,829, issued to Chen et al., U.S. Patent No. 5,742,343, issued to Haskell et al., and U.S. Patent No. 5,619,256, issued to Haskell et al., are concerned with digital compression techniques. Specific attention to compression of digitized video is found in U.S. Patent No. 5,764,803, issued to Jacquin et al. Compression of 3-dimensional images is treated by U.S. Patent No. 5,612,735, 10 issued to Haskell et al.

The evolving techniques of digitized video transmission have resulted in a transmission standard, called Asynchronous Transfer Mode (ATM). U.S. Patent No. 5,668,841, issued to Haskell et al., describes data transmission using the ATM standard. An ATM converter is discussed in U.S. Patent No. 5,809,022, issued to Byers et al. U.S. Patent No. 5,724,349, issued 15 to Cloonan et al., suggests an approach to packet switching within an ATM system. An ATM architecture is discussed in U.S. Patent No. 5,781,320, issued to Byers. Interfacing to ATM systems is addressed in U.S. Patent No. 5,842,111, issued to Byers.

A solution to the PPV problems noted above utilizing digitized video has been termed, Video on Demand (or VOD). In a VOD system, digitized video programming is made available 20 to individual cable television subscribers in response to specific requests made by the user. U.S. Patent No. 5,867,155, issued to Williams, describes the use of VOD for a very specialized

application. Sea Change, International, has proposed a VOD approach for cable television subscribers. U.S. Patent No. 5,583,561, issued to Baker et al., assigned to the assignee of the present invention and incorporated herein by reference, discloses and teaches a complete, modern VOD system employing a centralized architecture utilizing an enterprise server developed by

5 Unisys Corporation.

To fully utilize VOD for commercially useful subscriber levels, substantial administrative message traffic is required between the subscribers and the central controlling site(s). For current PPV, this typically takes the form of telephone communication with the subscriber. This is acceptable for PPV only because the total programming selections and timing are so limited.

10 However, VOD on demand so increases the options available to the subscriber, such manual attention is impractical.

To efficiently utilize the resources of the VOD system requires a methodology for transferring programming materials (or assets) from long term storage (i.e., disk) to an appropriate memory space and initiating streaming from an appropriate video server. When an asset is requested which is already being streamed, no additional transfer is required, if the associated video server has sufficient capacity to also stream the asset for the new request.

SUMMARY OF THE INVENTION

The present invention overcomes many of the disadvantages found within the prior art by

5 providing a video on demand system which separates the tasks of supplying video to subscribers from the tasks associated with managing the subscriber interface. The subscriber message interfaces can thus be automated such that most administrative matters are easily handled without operator attention. A message protocol is provided which enables the system to properly sequence and manage the message traffic.

10 The key to this approach is to provide an architecture in which the necessary functions are divided into two separate portions. A first hardware and software subsystem, called a video server, is specifically dedicated to retrieving and transmitting the stream of video information. Virtually no other functions are performed by the video server. A second hardware and software subsystem, called the transaction server, handles virtually all other functions including control 15 interface with the subscribers, digitized video data storage, subscriber accounting, etc. In accordance with the present invention, the transaction server automatically manages message traffic with the subscribers using the preferred protocol.

The video server has two primary design criteria. First, it must be highly optimized to handle the extremely high input/output data rates. In essence, this is the sole function of the video 20 server, and therefore, the design of the video server hardware and software are most directed towards this characteristic. Because the role of the video subsystem is simplified and single

dimensional, video subsystems utilizing current technology can be produced at a surprisingly low cost.

The second major design criterion of the video server subsystem involves modularity. The addition of active subscribers, each viewing individual video programs (or the same program at different times), tends to increase the total input/output load of the video server subsystem linearly. Therefore, there is great economic incentive to design the video server subsystem in such a manner that the hardware resources to implement the video subsystem may be linearly increased in relatively small (and inexpensive) increments. In the preferred mode of the present invention, the video server subsystem consists of one or more input/output data rate optimized, industry compatible computers operating under a readily available, commercial operating system, such as Windows NT. Using 3 mbytes/second per video stream as a standard, each such device can effectively service thousands of different and independent video streams. Within each video server, storage can be added to handle more video programs and communication interfaces can be added to provide more video streams. Therefore, the same design architecture and components are suitable for a wide range of system sizes, and the capital cost to the video programming supplier can be readily determined as further subscribers are added to the system.

Unlike the video server subsystem which is optimized to provide a low cost, highly modular approach to a single function, the transaction server is optimized to provide a low cost approach to a wide and highly expandable variety of functions. In fact, all of the functions of the video on demand system, except for the video streaming function performed by the video server, are accomplished by the transaction server. Typical tasks include: transactional interface with the

subscribers, subscriber fee accounting, requested program storage, video server subsystem control, receiving video from a satellite and storing it in an archive, etc. Thus, the ideal hardware/software platform for implementation of the transaction server is a readily expandable, highly flexible, high availability, highly recoverable, realtime oriented mainframe system. In the
5 preferred mode of the present invention, the Unisys 2200 is used to host the transaction server.

In accordance with the present invention, a subscribing user transfers a programming request to the transaction server. The transaction server makes the required subscriber accounting entry and notifies the corresponding preloaded video server platform of the new subscriber request. If the asset is not preloaded, in addition to the subscriber accounting, the
10 transaction server must access the requested video program and spool it to the video server. Depending upon the rate of use of the requested video program, the data might be stored in memory (for high volume use), on a disk or other mass storage device (for medium volume use), or in some other medium (for low volume use).

In the preferred mode of the present invention, the user is permitted to pause, reverse, or
15 fast forward the requested video program though commands entered from the set top subscriber box. These functions are intended to appear similar to normal VCR commands to the user. These commands are sent to the transaction server which utilizes them to control the corresponding video stream output of the video server subsystem. Thus the user is provided with all of the advantages of VCR rental without the need to physically transport the medium (i.e.,
20 cassette tape) back and forth between the rental store and the user site.

In view of the power and flexibility of the transaction server, other diverse but somewhat related functions may be provided. For example, a user might order a pizza delivery via the set top subscriber box to transaction server interface, or the user might access the internet, e-mail, or faxes via the transaction server. If this interface is implemented over a readily available, 5 publically accessible, network, such as the internet, many additional functions are possible.

As can be readily appreciated, the volume of message traffic between the subscribers and the transaction server can become substantial in order to implement these functions. Furthermore, it is necessary for the transaction server to manage these messages, because the set top subscriber box at the subscriber site is typically a rather "dumb" device having limited or no 10 programmability. The preferred embodiment of the present invention employs the basic UDP protocol for data transfers between the set top subscriber box and the transaction server. However, in order to properly manage the large number of messages, each is tagged with a sequence number and set top subscriber box address. These messages also contain a function code which identifies the basic content of the message and indicates whether a particular message 15 needs to be acknowledged.

In accordance with the preferred mode of the present invention, only one copy of a given program or asset exists on disk within the video server. Even when an asset is memory streaming a copy of the asset must still exist on disk. The methodology enables the video server to combine multiple physical disks into one or more logical volumes. In loading assets into a video server, the 20 low usage location is chosen which has acceptable free space having the least current percentage of performance utilization, PUT (PUT = StreamCount * AccessBitRate).

The video server memory is divided into blocks of 4MB and is configurable. When an asset is loaded into memory, it is allocated the number of blocks necessary to hold the asset. When an asset is unloaded the blocks are returned to the available pool. The blocks need not be contiguous to hold an asset. A special video server load disk message procedure provides to the 5 loading of an asset to a video server. This procedure describes when an asset stream request should use disk streaming and when an asset stream request should initiate memory streaming.

A video load/swap memory procedure determines how the video server controls the loading and swapping of memory. Only one memory load or swap may be active at a time.

The MAS/VS (Multimedia Application Server/Video Server) command/status protocol 10 controls the actual streaming functions. Streaming from the beginning of the asset (BOA) can begin whenever asset loading reaches the point (i.e., checkpoint) from which the streaming of data will never catch up with the loading function. Other streaming functions (i.e., streaming from other than the BOA must await complete loading of the entire asset. Streams are moved from disk to memory, as the memory asset is sufficiently loaded to allow the transition.

15 Assets are unloaded as required to make most efficient use of available memory.

Additional message traffic is necessary to notify the MAS of the asset status at each video server.

BRIEF DESCRIPTION OF THE DRAWINGS

Other objects of the present invention and many of the attendant advantages of the present
5 invention will be readily appreciated as the same becomes better understood by reference to the
following detailed description when considered in connection with the accompanying drawings, in
which like reference numerals designate like parts throughout the figures thereof and wherein:

Fig. 1 is a schematic diagram showing the operation of the overall video on demand
system of the present invention;

10 **Fig. 2** is a schematic diagram showing spooling of the video programming data from
typical mass storage devices;

Fig. 3 is a schematic diagram showing the generation of a video stream from spooled data
within a memory subsystem;

15 **Fig. 4** is a schematic diagram showing video streaming as synchronized on one minute
boundaries;

Fig. 5 is a schematic diagram showing operation of a video server platform;

Fig. 6 is a schematic diagram showing video streaming of multiple programs from a single
video server platform;

20 **Fig. 7** is a schematic diagram showing video streaming from video programming data
spooled on disk drive mass storage units;

Fig. 8 is a schematic diagram showing video streaming from video programming data stored on both disk drive mass storage units and memory subsystems;

Fig. 9 is a block diagram of a maximum configuration video server;

Fig. 10 is a detailed diagram of the operation of the transaction server of the preferred

5 mode of the present invention;

Fig. 11 is a system diagram of the VOD network showing the major data and message paths;

Fig. 12 is a system diagram showing a high reliability system employing redundancy;

Fig. 13 is a system diagram showing major message paths;

Fig. 14 is a diagram showing the basic message format;

Fig. 15 is a table of the defined function codes;

Fig. 16 is table showing definitions of various control terminology;

Fig. 17 is a flow diagram showing operation of the MAS/VS load-disk message procedure;

Fig. 18 is a flow diagram showing operation of the MAS/VS start-streaming message procedure;

Fig. 19 is a flow diagram showing operation of the MAS/VS load-swap procedure; and

Fig. 20 is a flow diagram showing operation of the MAS/VS selection procedure.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Fig. 1 is a schematic diagram 10 showing the overall video on demand system of the present invention. A subscribing user (not shown) is positioned adjacent standard television receiver 34. Through this television receiver, the user is capable of viewing video programming material transferred to his location via coaxial cable 30 from network 26 in the fashion currently known in the cable television industry. The interface between coaxial cable 30 and standard television receiver 34 is provided by set top subscriber box 32, which provides the conversion between MPEG-2 digitized video format and the analog video signal required by television receiver 34.

In many respects, set top subscriber box 32 is similar to the set top subscriber boxes utilized with existing cable television systems with the slight functional differences described in more detail below. The basic reason for these slight differences is to permit a subscribing user to communicate with transaction server 12 in a two directional manner. Not only does set top subscriber box 32 receive video programming data via coaxial cable 30 and present it to television receiver 34, but set top subscriber box 32 is capable of transferring user requests via coaxial cable 30 and network 26 to transaction server 12 via path 28. The most important requests in accordance with the present invention are those which initiate and control the individualized video on demand programming.

When the user is interested in viewing a particular video program, a request is made from set top subscriber box 32 and transferred to transaction server 12 via coaxial cable 30, network 26, and path 28. Transaction server 12, a Unisys 2200 system in the preferred embodiment, is

provided access to video programming information from satellite receiver 14, from analog video storage 16 and digital mass storage 18. In each instance, the video programming data is either received in digital form or converted to digital form. According to the preferred embodiment of the present invention, the MPEG-2 standardized format is utilized.

5 Whenever a request is received, transaction server 12 checks various security parameters, makes appropriate subscriber billing entries, and generally performs all of the necessary administrative functions as described below in greater detail. Additionally, transaction server 12 stores digital video data for transmission by the video server assigned to the requesting subscriber. One of video server platforms 20, 22,, or 24 is assigned the task by transaction server 12 and 10 the stored digital video data is supplied via the digital data bus shown. In the preferred mode of the present invention, each video server platform is a separate industry compatible, Windows NT based, computer platform. Once transferred to the selected video server, the requested video programming is transmitted via network 26 and coaxial cable 30 to set top subscriber box 32 and television receiver 34.

15 **Fig. 2** is a schematic diagram showing the spooling of data from digital disk mass storage devices. For the preferred mode of the present invention, the digitized video programming data is stored in MPEG-2 format. In the spooling process, the MPEG-2 organized and placed into memory as a programming file 55. To produce commercially acceptable video, 3mbits/second is required. That means that a two hour video program requires the about 2.7 billion bytes of data 20 storage. Because of cost considerations, many of the programs having low and moderate usage will best be stored on mass storage disk until requested

Individual storage disks 48, 50, ..., and 52 each store a number of video programs in MPEG-2 format. As requested, this data is transferred via storage bus 46 through disk control 42 through I/O bus 38 and placed in memory 55 via path 40. A software program spools the data to the ATM interface 54 at the required speed.

5 **Fig. 3** is a schematic diagram showing spooling of high volume digitized video program. For those programs having a high user demand, it is much more efficient to spool the program files from random access memory rather than mass storage disk systems. In this context, high volume means a high probability that the given program will be in use during high service volume periods. That means that there will need to be random access storage allocated to the storage of 10 that given program during peak memory demand. As a result, the system should simply allocate random access storage to that given program. Very popular, recent movies are typical of such high volume programs.

If a program is a high volume program, it is preferably stored in auxiliary memory 56. Upon request, software residing in memory 36 directs the storing of data from auxiliary memory 15 56 and transferring it via path 40 and I/O bus 38 to ATM interface 54. It should be noted that this is significantly more efficient than the storing operation shown in Fig. 2, since the video data only needs to be read out of memory instead of having to be loaded from disk each time the data is used. Furthermore, there is no additional cost if a program is of sufficiently high volume that the required random access memory must be allocated to the program anyway.

20 **Fig. 4** is a schematic diagram 58 showing the synchronization of a given video program around one minute time slots. In concept, the present invention provides subscribers with video

on demand. However, as a practical matter, by synchronizing multiple users around one minute time slots, the maximum number of transmissions to all users of the given video program cannot exceed 60 per hour of programming and 120 for a two hour standard video program. That means that for a given high volume program (which may be requested by hundreds or even thousands
5 within the length of time to view the program) each requester is assigned to an appropriate time slot.

First time slot 60 provides the first minute of video programming to one or more requesters. During one minute time slot 62, the initial requesters receive the second minute of programming, and one or more requesters may be starting with the first minute of programming.
10

At the nth time slot 64, the initial requesters are viewing the nth minute of programming, the second group of requesters is viewing the n-1 minute of programming, and the nth group of requesters is viewing the first minute of programming. At final time slot 66, the initial requesters are viewing the final minute of programming, the second group of requesters is viewing the second to last minute of programming, and a new group of requesters is viewing the initial minute
15 of programming.

The reduction in total data requirements utilizing these one minute time slots is substantial. Commonly assigned U.S. Patent No. 5,583,561, issued to Baker et al., incorporated herein by reference, discusses this feature in greater detail. The total delay to a requester is no more than one minute and will average one half minute, making the process perfectly acceptable and barely
20 noticeable to the subscribers.

Fig. 5 is a schematic diagram of a single industry compatible, Windows NT based video server platform. The video server subsystem is composed of a number of separate and largely independent video server platforms. Each is configured to have a maximum memory configuration and maximum I/O configuration. Digitized video programming data in the MPEG-2 format are moved from transaction server 68 via interconnect 70 into the assigned video server platform. Program 74 and program 72 are shown. These programs are placed onto network 78 under control of transmission control software 80 for transfer to the requesting subscriber(s). For a given program being sent to a single user, transmission control software 80 simply retrieves the video data from memory in a sequential fashion at 3 mbytes/second and places it on network 78.

Fig. 6 is a schematic diagram showing transfer of high volume program 82 into the video server platform of Fig. 5. The transfer is performed by the transaction server as discussed above. The transferred data is transferred to the video server platform via I/O bus 84 Up to ten programs can be stored and streamed from a single video server For simplicity, only one video server is shown. For a view of multiple video servers within a system, refer to Fig. 1.. All other referenced elements are as previously described.

Fig. 7 is a schematic diagram showing the spooling of low to moderate volume digitized video program data. For lower volume programs, storage on disk storage mass memory may be appropriate. A low volume video program is one in which it is highly unlikely that more than one request is received during the runtime of the video program. Therefore, the servicing of the request is most probably an index sequential task for retrieving the data and transmitting it to the user. This is readily distinguishable from the high volume video programs for which transmissions

within multiple and perhaps many of the one minute time slots is expected (see also Fig. 4).

These programs are spooled to the video server platform as shown. The remaining referenced elements are as previously described.

Fig. 8 is a schematic diagram showing transfer of low and high volume video programs to
5 the same video server platform. All referenced elements are as previously described.

Fig. 9 is a block diagram 84 of the maximum configuration of the transaction server of the
preferred mode of the present invention. In this preferred mode, the video server is implemented
using a current model Unisys mainframe system. In accordance with this product, the system is
expandable from a single processor, single main memory, and single I/O controller to the
10 maximum system shown.

Instruction processors 102, 104, 106, 108, 110, 112, 114, and 116 communicate with
main memories 86, 88, 90, and 92 via crossbar interconnects 94, 96, 98, and 100. Each
instruction processor may be coupled with up to four third-level caches, as shown. Direct
Input/Output bridges 118, 120, 122, 124, 126, 128, 130, and 132 each handle video output
functions. Each of the direct Input/Output bridges may be partitioned into separate partitions as
15 shown. Additional description of partitioning may be found in U.S. Patent Application Serial No.
08/779,472, filed January 7, 1997, commonly assigned to the present invention and incorporated
herein by reference.

Fig. 10 is a detailed functional diagram 134 of the transaction server. Communication
20 with set top subscriber box 32 (see also Fig. 1) is managed by set top management module 142.
Initial requests are selected by the user and honored through utilization of menu transaction

module 140. After initiation of the servicing of a given request, control of the matter is given to session manager 138 for completion. Any and all communication with the transaction server are monitored by security module 148. Administration module 136 provides overall control of the transaction server.

5 The transaction server may be utilized to interface with the internet. The selected hardware and software system selected for the preferred mode provide internet server facilities in a commercially usable form. Video server session management module 146 provides the detailed functions (e.g., spooling of digital video programming) associated with the primary video on demand service. These control functions are directly interfaced to the video server subsystem via
10 video server interface 150.

Event logging module 154 journals the functions performed. This log is made available to digital network control services 152. Media directory services and asset management module 162 provides long term control and asset management. Historical storage of these data is performed by asset storage management and asset capture.

0 15 In performing the actual video on demand service, the appropriate requested digitized video program is accessed from databases 160. It is spooled by asset delivery, video streaming module 158. The transfer is made via video server interface 156 (see also Fig. 1).

Fig. 11 is a diagram 170 of the overall VOD system with particular emphasis upon the data and message transfer paths. Multimedia Application Server (MAS) 172 accesses Resource
20 and Subscriber Database 186 via interface 174 and Asset Database 188 via interface 176. Similarly, Asset Archive Storage 190 is accessed via interface 180.

Multimedia Application Server 172 interfaces with all other major system elements via path 182. ATM channel adapter 192 provides the link using the ATM protocol to interface 194 and ATM network 200.

ATM interfaces include path 196 to router 198 which communicates with workstations

- 5 204 via radio frequency path 202. Router 212, coupled via path 208, also interfaces to Local Area Network (LAN) and to workstations 218 and 220 via stubs 214 and 216, respectively.

ATM 200 is also coupled directly to individual set top 228 via path 210 and to VODA 222 for streaming video along path 224 to QAM modulator 230, which supplies video to set top 228.

Fig. 12 is a diagram showing a highly reliable video on demand system. The high

- 10 reliability is achieved by having Multimedia Application Server (MAS) 232 completely duplicated as Multimedia Application Server 266. During peak loading, both can provide subscriber services. During periods of lesser loading, one or the other may be removed from service due to failure, maintenance, etc. XPC file accelerator 258 couples these two duplicate entities. Each Multimedia Application Server (i.e., 232 and 266) are redundantly coupled to Resource and
15 Subscriber Database 260, Asset Database 262, and Asset Archive Storage 264.

Within Multimedia Application Server 232 (and corresponding MAS 266) are found the Multimedia Application Server core software 240 (and corresponding core software 274) which communicates via middleware 242 (and middleware 276). Communication is with Relational Database 234 (268) via path 238 (272), and Message Retention Services 2 256 (280) via path 254 (278). CPCCOMM 252 (282) and ATP 250 (286) communicate directly with Video Server Frame

and Stream Spooling 236 (270) via path 248 (284). Similarly, communication is had with FTP 246 (290) via path 244 (288).

Fig. 13 is diagram 300 which conceptually shows the major message paths in a Sonnet Ring configured system. Multimedia Application Server 310 directly communicates with system 5 components: Asset Storage 308, Asset Management 306, ISP Connection 302 and Filling System 304.

Multimedia Application Server 310 receives messages directly from set top subscriber box 322 via path 312. It transfers messages to set top subscriber box 322 through path 314 to Sonnet Ring 316 to Hub 318 and then through cable 320. This path also provides the video for display 10 on television receiver 324.

Fig. 14 is a diagram 326 of the message format used to communicate between Multimedia Application Server 310 and set top subscriber box 322 (see also Fig. 13). Field 328 contains a four byte integer giving the version level. The current version is 2.

Field 330 is designated Set Top Command and Control Protocol (STCCPL). It is a four 15 byte integer which specifies the length in bytes of variable length field 242. Field 332 is a four byte integer which provides the sequence number. This sequence number permits the Multimedia Application Server and set top subscriber box to synchronize message traffic by establishing the order in which messages have been generated. The sequence number is created by incrementing the sequence number of the previous message. For those messages which must be acknowledged 20 (see below), the sequence number permits coordinating messages and acknowledgments.

The basic function code is specified in the four byte integer of field 334. These function codes are defined in more detail below. The Multimedia Application Server defines a unique four byte integer which is inserted into field 336. This is defined as the Connection ID. Field 338 provides for the entry of eight one byte integers. In the preferred embodiment, only the first five 5 integers are used to uniquely address one set top subscriber box within a network. The remaining three one byte integers are left blank.

Field 340 is the packet type. It is specified as a four byte integer. The packet type is undefined for messages sent from a set top subscriber box to the Multimedia Application Server. For every message the Multimedia Application Server receives, it sends a response back to the set 10 top subscriber box containing the STCCP output. This packet type is given the value of 1. An unsolicited administrative message sent from the Multimedia Application Server to the set top subscriber box which does not require an acknowledgment is defined as packet type 2. A similar message requiring an acknowledgment is defined as packet type 3..

Variable length field 342 may be from 0 to 952 bytes in length. The length is specified by 15 field 330. This provides great flexibility in creating new and unique messages for specific occasions.

Fig. 15 is a table showing the definition of the basic function codes to be entered into field 334 (see also Fig. 14). The function code of 1 is used to indicate to the Multimedia Application Server that a particular set top subscriber box has been switched on.

20 Messages sent from the Multimedia Application Server to one or more set top subscriber boxes are assigned the function code of 2. A message having a function code of 3 notifies the

Multimedia Application Server that the corresponding set top subscriber box has been powered down.

The set top subscriber box acknowledges messages from the Multimedia Application Server using a function code of 4. Reinitialization of a set top subscriber box is identified by a 5 function code of 5.

Fig. 16 is table showing the definitions of a number of significant terms associated with the management of asset transfer and streaming. The acronyms in the left column are defined by the immediately opposite entry in the right column.

Fig. 17 is a flow diagram showing the operation of the MAS/VS (Multimedia Application Server/Video Server) load-disk message procedure. The load-disk message is sent from the MAS to the video server to request that an asset be loaded. The message input parameters are AssetID (identifies the program asset to be loaded) and AssetSize (identifies the memory space required).

Entry is via element 350. At element 352 the ID and size parameters are assembled. The message is transferred from the MAS to the video server at element 354. The two basic 15 responses are that there is sufficient space or there is not sufficient space. In the latter case, element 358 sends control to element 374 for the sending of an error message and exit via element 380. This situation corresponds to insufficient system resources to accommodate the current user request.

Assuming the sufficient space exists, element 358 gives control to element 360 to ensure 20 that the streaming status is cleared. Element 362 establishes the identity of the asset to be loaded. The performance utilization variable is cleared by element 364. Element 366 sends a message

notifying MAS that the loading operation is in process, and actual data transfer is initiated by element 368.

Element 370 determines when the transfer is complete. If yes, control is given to element 356 to send the status message and to element 380 for exit. If not, element 372 determines 5 whether a transfer error has occurred. If yes, element 374 sends an error message and transfers control to element 380 for exit. Element 376 determines whether a sufficient portion of the subject asset has been transferred to begin streaming from the beginning of asset (BOA). This corresponds to the checkpoint. If the checkpoint has been reached, element 378 sends the message to MAS. If not, control is returned to element 370.

10 **Fig. 18** is a flow diagram showing operation of the MAS/VS start-streaming message procedure. This procedure begins streaming the program asset to the requesting user. The two input parameters are AssetID and AssetStart (which defines from where the streaming is to begin because the streaming need not begin from BOA).

15 The procedure is entered at element 382. Element 384 searches for the requested program asset. If the requested program asset is not located, element 386 transfers control to element 414 to send an error status message and exit via element 416. If the asset is located, element 386 gives control to element 388 which determines whether the requested asset is currently memory streaming. If no, element 394 determines if disk threshold is set. If yes, element 396 initiates a load/swap procedure as further explained in Fig. 19. Control is given to 20 element 400 to determine if there is insufficient disk capacity. If yes, element 410 sends a no-resources message and exit is via element 416.

Element 406 determines whether the asset has been loaded to the check point (i.e., sufficient loading to initiate streaming at BOA). If no, an error has occurred and control is given to element 414 to send an error message and exit via element 416. If the checkpoint has been reached, element 406 gives control to element 408 to begin streaming from BOA. Element 412
5 sends message acknowledging initiation of streaming, and exit is via element 416.

If element 388 does find the memstreaming flag set, control is given to element 408 to begin streaming. A message is sent by element 412 to indicate that the memory streaming has been initiated. Following transmissions of the status message, the procedure is provide an exit via element 416.

10 Fig. 19 is flow diagram showing operation of the MAS/VS memory load/swap procedure. This procedure is initiated every time a memory asset stream terminates or a start-stream message is received for a disk asset on a volume that is over threshold.

Entry is via element 418. Element 420 searches for the most active asset not currently streaming. If no asset is located (i.e., all loaded assets are currently in use) element 422 gives
15 control to element 480 for exit. No swapping or further loading is possible.

Element 424 determines whether the available free space is larger than the asset to be loaded. If not, element 426 searches for a potential asset to swap out to provide sufficient loading space. The search criteria include lower performance utilization than asset to be loaded and size of free space sufficient if found asset is to be swapped out of memory. If a suitable asset is found,
20 control is given to element 434 for unloading of the selected asset and the procedure is repeated, now having sufficient free space.

If no asset is found meeting the criteria of element 428, control is given to element 430 to search for any currently inactive assets having lower performance utilization. If no such assets are found, control is given to element 480 for exit. If any such assets are found, control is given to element 434 for unloading of the selected asset and the procedure is repeated.

5 When sufficient free memory has been secured, element 436 is given control for sending a memory load request memory. Element 438 waits until a favorable response is received. When a favorable response has been received, element 462 initiates loading of the memory. The file is loaded at element 464.

During the loading process, element 466 determines whether an error has occurred. If 10 yes, the loading process is terminated. Element 472 deletes the asset identifier from the memory directory, and element 478 sends an error message. Exit is via element 480.

Element 468 determines when the transfer operation has been completed. When finished, control is given to element 474 to send a completion message, and exit is via element 480. The loading process is monitored by element 470. When the checkpoint is reached, element 476 sets 15 the associated flag. This means that a sufficient portion of the asset has been loaded that streaming can begin from BOA (i.e., beginning of asset). Control is returned to element 466.

Fig. 20 is a flow diagram showing operation of the MAS/VS selection procedure. This procedure basically assigns user requests to available video servers. It is implemented through the use of a MAS database containing data on:

20 What assets are loaded on each video server;

What assets are memory streaming on each video server; and

The remaining disk capacity on each video server.

This data is maintained in response to the status messages discussed above and below.

The procedure is entered at element 482. Element 484 searches all servers for an available video server with the subject asset already loaded. If element 486 finds the video server, control
5 is given to element 488, which determines if the asset has been completely loaded. If yes, control is given to element 518 for initiation of streaming. This situation corresponds to a new user request for a program which is already in use.

If element 488 determines that the asset is not completely loaded, control is given to element 490, which determines if the asset has been loaded to the check point and the request
10 involves streaming from the BOA. If yes, control is given to element 518 to begin streaming. If element 490 conditions are not met, it is because the loading has not reached check point or the user request does not involve streaming from the BOA. Element 492 deals with the BOA situation. If yes, control is given to element 490 to await loading to the check point. If no, control is given to element 494 for return to element 486 to await complete loading.

15 If no video server is found meeting the criteria of element 486, control is given to 496 to continue to search for an appropriate available video server and wait until one is available. If and when the appropriate video server is found, control is given to element 502 to determine when to initiate streaming. Element 502 determines whether the asset is completely loaded. If completely loaded, control is given to element 518 to initiate streaming. If asset not completely loaded,
20 element 512 determines whether check point has been reached and request is for streaming from BOA. If yes, control is given to element 518 to initiate streaming.

If element 514 finds that a BOA start has been requested, control is given to element 520 to wait for loading to reach the check point. Otherwise, control is given to element 522 for completion of the loading process.

When element 498 finds no suitable video server, control is given to element 500 for any available video server into which the requested asset may be loaded. If no video server is found by element 504, control is given to element 536 for exit. If the video server is found, control is given to element 506 which initiates loading of the requested asset. Loading errors are determined by element 508. If found, the error is logged by element 510 and control returns to element 484 to restart the procedure. If no video server error is found, element 516 determines video server status. If yes, control is returned to element 484 for restarting.

Otherwise, streaming is initiated at element 518. Element 524 determines whether resource limited. If yes, element 526 issues start streaming message. Element 528 searches for a video server error, which is logged by element 530. If streaming has not started, element 534 indicates that the request was rejected. Exit is via element 536.

Having thus describe the preferred embodiments in detail, those of skill in the art will be readily able to use the teachings found herein to make and use yet other embodiments within the scope of the claims appended hereto.

What is claimed is: